

1 Thank you for your detailed and helpful reviews! Reviewers 2 and 3 are positive about the paper and don't have specific
2 questions, so we will focus on responding to Reviewer 1.

3 Reviewer 1's main concern seems to be the applicability of the model, pointing out that the example in the introduction
4 is "toy," and being concerned about the NP-hardness of the problem in the fully general case, saying she/he appre-
5 ciates Section 5 where everything can be done efficiently, but wondering whether we have done enough in terms of
6 characterizing easy cases. Let us address these concerns next.

7 First, the purpose of the example in the introduction is more to illustrate the model (and Reviewer 3 liked the example
8 for that reason) than to illustrate a direct practical application. The most obvious applications are in the special case
9 where the agent has some data and can withhold some of these data, or even parts of individual examples. (Indeed, this
10 is the case that motivated us to study this topic in the first place; but in the end, we had results for a much more general
11 model so we decided to present it accordingly.)

- 12 • For example, we can think about someone applying to a music school and submitting recordings of her
13 performances — but she could easily omit some recordings, or cut parts of recordings out. There are many
14 similar examples involving not musicians but athletes, actors, etc.
- 15 • Another example that we took out of the paper due to space constraint is a publisher who is trying to convince
16 an advertiser that high-value users visit her site during a trial period. In this case, the publisher can decide to
17 not show low-value users the ad, or direct them to a different part of the site.

18 These examples illustrate the kinds of settings for which our graph-theoretic framework makes more sense than the
19 costly feature manipulation framework. E.g., conceptually, for the example of cutting off the end and beginning of a
20 recording, the graph would have edges from intervals (of the recording) to all subintervals. For the example where the
21 publisher shields the advertiser from certain users, the graph would have edges from every user (sample) to both itself
22 and the "empty" sample/signal (no user is shown).

23 These examples, and indeed all the settings that originally motivated this work, also exactly fit the special case in
24 Section 5, so everything is tractable too. Section 5 captures a very general class of applications, which still involves a
25 misreporting graph. While it is true that there may be other special cases that are also tractable, in general, there is no
26 good way to characterize all special cases that are not NP-hard. In terms of economic theory, there is a very natural
27 distinction between Section 5, where the key aspect is that the revelation principle holds (i.e., without loss of generality,
28 good types can report truthfully) and the more general case where the revelation principle does not hold. (In other,
29 non-learning contexts in mechanism design, whether the revelation principle holds is often an aspect that determines
30 whether things are tractable.)

31 We will add these examples to the paper; we agree that they should be helpful to the reader.

32 We thank all the reviewers again for their detailed and encouraging comments!